The Invisible Hand of Dynamic Market Pricing

VINCENT COHEN-ADDAD, Ecole normale supérieure, Paris, France, vcohen@di.ens.fr
ALON EDEN, Tel-Aviv University, Israel, alonarden@gmail.com
MICHAL FELDMAN, Tel Aviv University and Microsoft Research, Israel, michal.feldman@cs.tau.ac.il
AMOS FIAT, Tel-Aviv University, Israel, fiat@tau.ac.il

Walrasian prices, if they exist, have the property that one can assign every buyer some bundle in her demand set, such that the resulting assignment will maximize social welfare. Unfortunately, this assumes carefully breaking ties amongst different bundles in the buyer demand set. Presumably, the shopkeeper cleverly convinces the buyer to break ties in a manner consistent with maximizing social welfare. Lacking such a shopkeeper, if buyers arrive sequentially and simply choose some arbitrary bundle in their demand set, the social welfare may be arbitrarily bad. In the context of matching markets, we show how to compute dynamic prices, based upon the current inventory, that guarantee that social welfare is maximized. Such prices are set without knowing the identity of the next buyer to arrive. We also show that this is impossible in general (e.g., for coverage valuations), but consider other scenarios where this can be done. We further extend our results to Bayesian and bounded rationality models.

1. INTRODUCTION

A remarkable property of Walrasian pricing is that it is possible to match buyers to bundles, such that every buyer gets a bundle in her demand set (i.e., a set of items S maximizing $v_i(S) - \sum_{j \in S} p_j$), and the resulting allocation maximizes the social welfare, $\sum_i v_i(S_i)$ (S_i being the bundle allocated to buyer i). However, Walrasian prices cannot coordinate the market alone; it is critical that ties be broken appropriately, in a coordinated fashion.

Consider the following scenario: two items, a and b, and two unit demand buyers, Alice and Bob. Alice has value $R \gg 1$ for item a and value one for item b, Bob has value one for each of the two items a and b. There are many Walrasian pricings in this setting, for example a price of $R-1$ for item a and 0 for item b. Indeed, assigning item a to Alice and item b to Bob under these prices maximizes simultaneously the individual utility of each buyer and the social welfare.

However, in real markets, buyers often arrive sequentially, in some unknown order, and get no guidance as to how to break ties. For these prices, ($p(a) = R - 1$ and $p(b) = 0$), if Bob arrives first then he will indeed choose item b, leaving item a for Alice to purchase, resulting in a social welfare maximizing allocation. If, however, Alice arrives first, she has equal utility ($= 1$) for both a and b and may select item b, so Bob will walk away without purchasing any item, which results in social welfare 1, compared with the optimal social welfare of $R + 1$. We furthermore remark that setting prices of

---

The work of V. Cohen-Addad was partially supported by ANR RDAM. The work of M. Feldman and A. Eden was partially supported by the European Research Council under the European Union's Seventh Framework Program (FP7/2007-2013) / ERC grant agreement number 337122. The work of A. Eden and A. Fiat was partially supported by ISF 1841/14. This work was done in part while A. Fiat and M. Feldman were visiting the Simons Institute for the Theory of Computing.

EC '16, July 24-28, 2016, Maastricht, Netherlands

http://dx.doi.org/10.1145/2940716.2940730

$p(a) = R$ and $p(b) = 1$, which are also Walrasian prices, could result in both Alice and Bob walking away, and resulting in zero social welfare.

One may suspect that we choose the wrong Walrasian pricing. It is known that in matching markets the minimal Walrasian prices coincide with VCG payments [Leonard 1983], and they are also the outcome of natural ascending auctions for matching markets [Crawford and Knoer 1981]. In this example the minimal Walrasian prices are to charge zero for both item a and item b. Indeed, if Alice arrives first, she will choose item a, and when Bob arrives he will choose item b, and this is the social welfare maximizing allocation. However, if Bob arrives first, he will be indifferent between the two items and may choose item a — again — this achieves a social welfare of 2 compared with the optimal social welfare of $R + 1$. In fact, one can show that minimal Walrasian prices always induce ties in demand [Hsu et al. 2016]. Moreover, there exist markets that admit unique Walrasian prices, yet may achieve zero welfare. For example, consider a single item valued at 1 by both Alice and Bob. The unique Walrasian price is 1, which may result in both buyers walking away without purchasing the item.

In fact, we can show that no static prices (and thus no Walrasian prices) can give more than 2/3 of the social welfare for buyers that arrive sequentially. Consider unit demand buyers Alice, Bob, and Carl, and items a, b, and c. Alice values a and b at one, and has zero value for c, symmetrically, Bob values b and c at one and a at zero, and Carl values c and a at one, and b at zero. A two line proof shows that no static pricing scheme, $p(a)$, $p(b)$, and $p(c)$ can achieve more than 2/3 of the optimal social welfare. Assume all prices are strictly less than one, and assume, without loss of generality, that $p(a) \geq p(b) \geq p(c)$. Now, Alice arrives and chooses item b, Carl arrives and chooses item c, and finally Bob arrives — but there are no items left for which Bob has a non zero valuation. Note that if $p(a) \geq 1$ then item a will not be sold as whomever is to buy it may decide simply to walk away, the same holds for items b and c so assuming that all prices are strictly less than one holds without loss of generality, given that one assumes that the prices achieve $\geq 2/3$ of the optimal social welfare.

However, consider the following twist, which changes the prices after the first buyer arrives. In the scenario above, when Alice arrives first and chooses (without loss of generality) item a, change the prices so that Bob will choose b and Carl will choose c. This is easily done by setting new prices $p'(b) < p'(c)$. Irrespective of whomever arrives after Alice, the prices will ensure that all items get sold and social welfare be maximized.

Obtaining optimal social welfare is trivial via dynamic pricing if the pricing mechanism knew which buyer was to arrive next. The dynamic pricing mechanism could make use of infinite prices to reduce the choices available to incoming buyer so that only a bundle consistent with optimal social welfare can be selected. The key difficulty arises because the prices need to be set *before* the preferences of the next buyer to arrive are known.

Thus, this paper studies the issues of static and dynamic pricing for sequentially arriving buyers. Our main result is the following:

Main Theorem: For any matching market (*i.e.*, unit demand valuations), we give a poly-time dynamic pricing scheme that achieves the optimal social welfare, for any arrival order and irrespective of any tie breaking chosen by the buyers. We give a combinatorial algorithm to compute such prices[1].

[1]We are grateful to an anonymous reviewer who pointed out that this result can also be obtained via an LP formulation of the problem and using LP-duality.

We show that the existence of Walrasian prices does not, by itself, imply that there exist dynamic pricing schemes that optimize social welfare. In particular, we give an example (Section 4) of a market with coverage valuations (a strict subclass of submodular valuations), which has a unique optimal solution, and where Walrasian prices do exist, and yet no pricing scheme (static or dynamic) can get the optimal social welfare.

We offer some remedies for this impossibility result.

— We show that a market with gross substitutes valuations that has a unique optimal allocation always admits a *static item pricing* scheme that achieves the optimal welfare (Section 6).
— Moreover, while full efficiency is in general impossible, we argue that for *any* profile of valuations, there exists a static pricing scheme that achieves at least a half of the optimal social welfare. This result can be viewed as a generalization of the Combinatorial Walrasian Equilibrium of [Feldman et al. 2013]. In fact we adapt the static bundle prices computed in [Feldman et al. 2015] for Bayesian agents to achieve the one half guarantee of the optimal social welfare, for any class of valuations.
— We identify additional classes of valuations that admit dynamic pricing schemes that obtain the optimal social welfare: (1) where buyer i seeks up to k_i items, and valuations depend on the item, and (2) for superadditive valuations.

The following remark is in order. Gross substitutes valuations are known to be the frontier for the guaranteed existence of a Walrasian equilibrium [Gul and Stacchetti 1999]. They are also the frontier with respect to computational tractability [Nisan and Segal 2006]: one can compute the allocation that maximizes social welfare in polynomial time. *Are gross substitutes valuations also the frontier for achieving optimal welfare via a dynamic pricing scheme?* More formally:

Main Open Problem: Does any market with gross substitutes valuations have a dynamic pricing scheme that achieves optimal social welfare?

We generalize our results beyond the full information setting. We introduce a new mixed Bayesian/worst case model, where the population of buyers is chosen by a Bayesian process, and then the order of arrival is chosen by an adversary. The adversary chooses the next buyer to arrive, and knows both the realization of the population and the decisions made by the dynamic pricing algorithm thus far. In particular, any arbitrary Bayesian process for determining the order of arrival is just a special case of this worst case adversarial setting.

We then consider a model of fuzzy valuations that captures both buyer bounded rationality and seller uncertainty: The buyers need not choose something in their demand, but can choose to [almost] maximize utility, to within a small error term. This models cases where the buyers themselves are unsure of their valuation, or are inconsistent over time. The sellers do not know buyer valuations precisely, but only to within some small error.

For both models above, we extend our techniques to give dynamic pricing that achieves social welfare close to optimum, up to the inherent uncertainty in the models. In fact, all of the above elements of uncertainty can be combined.

Finally, a note on genericity is in order. If one were to perturb agent valuations, a unique social optimum would emerge, and it follows from our results that a static pricing would suffice to obtain optimal welfare for sequential sales with gross substitute valuations. However, one may consider perturbations of valuations as implying alternative tie breaking rules. This approach might lead to arbitrarily low social welfare. As pointed out in Section 7, we can devise dynamic prices that achieve good social welfare even in such cases.

1.1. Related Work

This paper combines issues of online computation and markets.

Walrasian equilibrium, where prices are such that optimal social welfare is achieved, and the market clears, given appropriate tie-breaking of preferences in the demand set dates back to 1874 [Walras 1874]. The existence of Walrasian prices for matching markets and more generally for gross substitutes valuations appears in [Kelso Jr and Crawford 1982; Gul and Stacchetti 1999]. We give a definition of these valuations in Section 6. Competitive analysis of online matchings were first studied in [Karp et al. 1990] where a randomized $1 - 1/e$ approximation to the size of the maximal matching was given.

The use of bundle pricing for Combinatorial Walrasian Equilibria (and no envy amongst buyers), while achieving one half of the social welfare, was given in [Feldman et al. 2013]. The use of static item prices for buyers arriving via a Bayesian process, with XOS valuations, which also achieves 1/2 of the optimal social welfare, was given in [Feldman et al. 2015].

The performance of posted price mechanisms was also studied under the objective of maximizing revenue in Bayesian settings, where it was shown to extract a constant fraction of the optimal revenue for single item settings [Blumrosen and Holenstein 2008] as well as for matching markets [Chawla et al. 2007, 2010a,b].

Recently, [Hsu et al. 2016] addressed a related problem regarding the role of prices in market coordination (in the absence of a central coordination entity). They focused on minimal Walrasian prices, and showed that in markets where there are multiple copies of every item, and under a genericity assumption of agents' valuations, such prices result in a minimal over-demand, and thus guarantee high social welfare. In contrast to [Hsu et al. 2016], our objective is to achieve optimal welfare precisely, and even without a genericity condition or multiple copies of an item. As we show, in this case dynamic prices are necessary.

More generally, our paper presents an additional application to the dynamic pricing paradigm [Cohen et al. 2015], where dynamic prices are set on future selfish decisions so as to achieve some predefined goal. In particular, this has been done in the context of minimizing the costs of selfish metrical matchings, selfish metrical task systems, and the selfish k-server problem.

1.2. The Structure of this Paper

In Section 2 we describe several types of pricing schemes for sequential buyers, static and dynamic, item prices and bundle prices.

In Section 3 we give a dynamic pricing scheme that achieves optimal social welfare, irrespective of how agents break ties, and for any order of arrival. We include a running example to help in clarifying the concepts and algorithms involved.

In Section 4 we show that dynamic pricing schemes cannot achieve optimal social welfare even if all of the following hold simultaneously: (1) Walrasian prices exist, (2) The socially optimal allocation is unique, and (3) The valuation is a coverage valuation.

In Section 5 we argue that the ideas in [Feldman et al. 2013, 2015] allow us to compute static prices that achieve 1/2 of the optimal social welfare, for any order of arrival, and any valuation.

In Section 6 we show how to compute static item prices that achieve optimal social welfare for sequentially arriving buyers if the valuation class is gross substitutes and the optimal allocation is unique.

In Section 7 we extend our results beyond the full information setting. We introduce a Bayesian model and a fuzzy valuations model and extend our techniques to give dynamic pricing for these models.

In the full version we show how to compute static bundle prices that achieve optimal social welfare for sequentially arriving buyers in the following two scenarios: (1) if the valuation class is super additive, and (2) if every bidder i seeks up to k_i items, and the item values depend only on the item.

2. MODEL AND PRELIMINARIES

Our setting consists of a set I of m indivisible items and a set of n buyers that arrive sequentially in some arbitrary order.

Each buyer has a valuation function $v_i : 2^I \rightarrow \Re_{\geq 0}$ that indicates her value for every set of objects, and a buyer valuation profile is denoted by $\mathbf{v} = (v_1, \ldots, v_n)$. We assume valuations are monotone non-decreasing and normalized (i.e., $v_i(\emptyset) = 0$). We use $v_i(A|B) = v_i(A \cup B) - v_i(B)$ to denote the marginal value of bundle A *given* bundle B. An *allocation* is a vector of disjoint sets $\mathbf{x} = (x_1, \ldots, x_n)$, where x_i denotes the bundle associated with buyer $i \in [n]$ (note that it is not required that all items are allocated). The *social welfare* (SW) of an allocation $\mathbf{x}$ is $\mathsf{SW}(\mathbf{x}, \mathbf{v}) = \sum_{i=1}^{n} v_i(x_i)$, and the optimal welfare is denoted by OPT($\mathbf{v}$). When clear from context we omit $\mathbf{v}$ and write SW($\mathbf{x}$) and OPT for the social welfare and optimal welfare, respectively.

An *item pricing* is a function $\mathbf{p} : I \rightarrow \Re^{\geq 0}$ that assigns a price to every item. The price of item j is denoted by $p(j)$. Given an item pricing, the *utility* that buyer i derives from a set of items S is $u_i(S, \mathbf{p}) = v_i(S) - \sum_{j \in S} p(j)$. The *demand correspondence* $D_i(I, \mathbf{p})$ of buyer i contains the sets of objects that maximize buyer i's utility; i.e., $D_i(I, \mathbf{p}) = \mathrm{argmax}_{S \subseteq I} u_i(S, \mathbf{p})$.

A *bundle pricing* is a tuple $(\mathcal{B}, \mathbf{p})$, where $\mathcal{B} = \{B_1, \ldots, B_k\}$ is a partition of the items into bundles (where $\bigcup_i B_i = I$ and for every $i \neq j$, $B_i \cap B_j = \emptyset$), and $p : \mathcal{B} \rightarrow \Re^{\geq 0}$ is a function that assigns a price to every bundle in $\mathcal{B}$. The price of bundle B_j is denoted $p(B_j)$. Given a bundle pricing $(\mathcal{B}, \mathbf{p})$, the utility that buyer i derives from a set of bundles S is $u_i(S, \mathbf{p}) = v_i(S) - \sum_{B_j \in S} p(B_j)$. The *demand correspondence* $D_i(I, \mathbf{p})$ of buyer i contains the sets of bundles that maximize buyer i's utility; i.e., $D_i(I, \mathbf{p}) = \mathrm{argmax}_{S \subseteq I} u_i(S, \mathbf{p})$.

We consider several types of pricing schemes: *static item pricing*, *dynamic item pricing*, *static bundle pricing*, and *dynamic bundle pricing*.

In static pricing schemes, prices are assigned (to items or bundles) initially, and never change then. In contrast, in dynamic pricing schemes, new (item or bundle) pricing may be set before the next buyer arrives. Item pricing schemes assign prices to items, whereas bundle pricing schemes partition the items to bundles and assign prices to bundles that are elements of the partition. Thus, the four types of pricing schemes are described as follows.

Static Item Pricing Scheme:

(1) Item prices, $\mathbf{p}$, are determined once and for all.

(2) Buyers arrive in some arbitrary order, the next buyer to arrive chooses a bundle in her demand set among the items not already allocated (and pays the sum of the corresponding prices).

Static Bundling Pricing Scheme:

(1) Bundles, and their prices, $(\mathcal{B}, \mathbf{p})$, are determined once and for all.

(2) Buyers arrive in some arbitrary order, the next buyer to arrive chooses a set of bundles in her demand set among the bundles not already allocated (and pays the sum of the corresponding prices).

Dynamic Item Pricing Scheme:

— Before buyer $t = 1, \ldots, n$ arrives (and after buyer $t - 1$ departs, for $t > 1$):
  (1) Item prices, $\mathbf{p}_t$, are set (or reset) before buyer t arrives, prices are set for those items that have not been purchased yet.
  (2) When buyer t arrives she purchases a set of items S in her demand among the items not already allocated (and pays the sum of the corresponding prices according to $\mathbf{p}_t$).

Dynamic Bundle Pricing Scheme:

— Before buyer $t = 1, \ldots, n$ arrives (and after buyer $t - 1$ departs, for $t > 1$):
  (1) A partition into bundles and bundle prices, $(\mathcal{B}, \mathbf{p}_t)$, is determined for the items that have not been purchased yet.
  (2) When buyer t arrives she purchases a set of bundles S in her demand set among the bundles on sale (and pays the sum of the corresponding prices according to $\mathbf{p}_t$).

We say that a pricing scheme achieves optimal (respectively, α-approximate) social welfare if for any arrival order and any manner in which agents may break ties, the obtained social welfare is optimal (resp., at least an α fraction of the optimal welfare).

3. OPTIMAL DYNAMIC PRICING SCHEME FOR MATCHING MARKETS

In this section we consider matching markets. Every agent seeks one item, and may have different valuations for the different items. Whereas this setting admits Walrasian prices, such prices are not applicable to the setting where agents arrive sequentially, in an unknown order, and choose an arbitrary item in their demand set.

We now describe a dynamic item pricing scheme for matching markets that maximizes social welfare — the sum of buyer valuations for their allocated items is maximized. The process we consider is as follows:

— The valuations of the buyers are known.
— The buyers arrive in some arbitrary order unknown to the pricing scheme.
— Prices are posted, they may change after a buyer departs but cannot depend upon the next buyer.

[Running example] To illustrate the process, we consider a running example of a matching market, buyers Alice, Bob, Carl, and Dorothy, items a, b, c and d. The valuations are given in Figure 3(a), where squares represent buyers, circles represent items, and A, B, C, D stand for Alice, Bob, Carol and Dorothy. The minimal Walrasian pricing is $p(a) = 1, p(b) = 7, p(c) = 7, p(d) = 0$. Under the minimal Walrasian pricing, or any static pricing, unless ties are broken in a particular way, sequential arrival of buyers will not produce optimal social welfare (as shown in the full version of the paper).
An example of the use of dynamic pricing that follows from our dynamic pricing scheme is given in Figure 3. Every row represents a phase in the process, where a single buyer arrives. The LHS graph in every row represents the valuations of the remaining buyers and items, thick edges represent a maximum weight matching. The RHS graph represents the graph of edges, upon which prices are calculated by Algorithm Price-Items. Directed cycles of length 0 (if any) are represented by thick edges. The arriving buyers along with the items they pick are specified in the right column.

The input consists of the graph $G = (N, I, \mathbf{v})$. G is a complete bipartite weighted graph, where N is the set of agents, I is the set of items, and for every agent $a \in N$ and item $b \in I$, the weight of an edge $\langle a, b \rangle$ is the value that agent a gives item b, $v_a(b)$ ($v_a : I \to \Re_{\geq 0}$ is the valuation function for agent a).

Without loss of generality, one may assume that in G we have that $|I| \geq |N|$, otherwise, we add dummy vertices to the I with zero weight edges to the vertices of the N side until $|I| = |N|$. OPT is the weight of the maximum weight matching in G (alternatively, the optimal social welfare). Let $M \subseteq N \times I$ be some matching in G, we define $\mathsf{SW}(M) = \sum_{(a,b)\in M} v_a(b)$ to be a function that takes a matching and returns the social welfare (value) of the matching.

We now continue to describe the dynamic pricing scheme. At time $t \in 0, \ldots, |N|$ (after the t-th agent departs), we define the following:

— $M_t \subseteq N \times I$ is the partial matching consisting of [a subset] of the first t agents to arrive, and the item of their choice, amongst the items available for sale upon arrival. The size of M_t may be less than t. It can indeed be the case that not all buyers may be matched since their demand set may be empty when they arrive.
— $N_t \subseteq N$ and $I_t \subseteq I$ are the first t agents to arrive and the items matched to them in the matching M_t.
— $N_{>t} = N \setminus N_t$ and $I_{>t} = I \setminus I_t$ are the remaining agents (to arrive at some time $> t$) and the items remaining after the departure of the t-th agent. Define $G_{>t}$ to be the graph G where agents N_t and items I_t have been discarded. *I.e.*, $G_{>t} = (N_{>t}, I_{>t}, \mathbf{v})$.
— We define $p_{t+1} : I_{>t} \to \Re_{\geq 0}$ to be the prices set by the dynamic pricing scheme after the departure of agent t (but before the arrival of agent $t + 1$).

To compute the function p_{t+1} we first construct a so-called "relation graph", $R_{>t}$, and then perform various computations upon it. The vertices of the relation graph are all items yet unsold, $I_{>t}$, the edges and their weights are as follows:

(1) Compute $M_{>t} \subseteq N_{>t} \times I_{>t}$, a maximum weight matching of the graph $G_{>t}$ which matches **all** vertices of $I_{>t}$.[2] For every item $b \in I_{>t}$, let $v_{>t}(b)$ denote the value of item b to the agent matched to item b in the matching $M_{>t}$.

[2]Note that such a maximum weight matching exists because initially $|N| \leq |I|$, and since every agent takes at most one item, $|N_{>t}| \leq |I_{>t}|$ continues to hold. Since all edge weights are non-negative, and $G_{>t}$ is a

(2) The edges of $R_{>t}$, denoted by $E_{>t}$, are a clique on the vertices $I_{>t}$, and their weights $W_{>t} : E_{>t} \to \Re$ are computed as follows: Let $M_{>t}$ be a maximum weight matching of remaining items and agents as defined above. For every pair $(a, b) \in M_{>t}$, and for every $b' \in I_{>t} \setminus \{b\}$ create an edge $\langle b, b' \rangle$. The weight of the edge $\langle b, b' \rangle$,

$$W_{>t}(\langle b, b' \rangle) = v_a(b) - v_a(b').$$

> **[Running example]** The initial graph $G_{>0}$ of our running example is given in Figure 3(a), where a maximum weight matching $M_{>0}$ is indicated by thick edges. The graph $R_{>0}$ is given in Figure 3(b). For example, the weight of the edge $\langle a, b \rangle$ is $v_{Alice}(a) - v_{Alice}(b) = -6$.

We give the following structural property of $R_{>t}$:

LEMMA 3.1. *There are no directed cycles of negative weight in $R_{>t}$.*

Proof: Assume there exists a negative cycle of length ℓ. Assume the cycle is comprised of $\langle b_1, b_2 \rangle, \langle b_2, b_3 \rangle, \ldots, \langle b_{\ell-1}, b_\ell \rangle, \langle b_\ell, b_1 \rangle$. This cycle corresponds to a cycle of alternating edges in $G_{>t}$ $(b_1, a_1)(a_1, b_2), (b_2, a_2) \ldots (a_{\ell-1}, b_\ell), (b_\ell, a_\ell), (a_\ell, b_1)$, where for every $j \in \{1, \ldots, \ell\}$, $(b_j, a_j) \in M_t$ and $(a_j, b_{j+1}) \notin M_t$.

For ease of notation, we define $\ell + 1 = 1$. According to the definition of weights in $R_{>t}$, we know that

$$\sum_{j=1}^{\ell} W_{>t}(\langle b_j, b_{j+1} \rangle) = \sum_{j=1}^{\ell} \left(v_{a_j}(b_j) - v_{a_j}(b_{j+1}) \right) < 0,$$

and therefore, $\sum_{j=1}^{\ell} v_{a_j}(b_{j+1}) > \sum_{j=1}^{\ell} v_{a_j}(b_j)$. We get that the matching M', which is constructed by removing the set $\{(b_j, a_j)\}_{j \in 1,\ldots,\ell}$ from $M_{>t}$ and adding the set $\{(b_{j+1}, a_j)\}_{j \in 1,\ldots,\ell}$, is of larger weight, in contradiction to $M_{>t}$ being a maximum weight matching.

■

We now process the relation graph $R_{>t}$:

(1) Let Δ be the smallest total weight of a cycle with strictly positive total weight in $R_{>t}$, and let $\epsilon = \frac{\Delta}{|I_{>t}|+1}$. Mark all edges in $E_{>t}$ that take part in **some** directed cycle of weight 0 in $R_{>t}$. Delete all marked edges. For every remaining edge e, set $W'_{>t}(e) = W_{>t}(e) - \epsilon$. Let $R'_{>t} = (I_{>t}, E'_{>t}, W'_{>t})$ be the resulting graph.
(2) Find a solution to the set of equations in Figure 1 by running algorithm Price-Items (see Figure 2) with $R'_{>t}$ as the input graph. Set $p_{t+1} = p$ where p is the output of Price-Items.

To show that indeed, $R'_{>t}$ can be used as an input for Price-Items, we show the following:

LEMMA 3.2. *All the directed cycles in ${R'}_{>t}$ are strictly positive.*

Proof: Let $\widetilde{R}$ be the graph which is obtained from $R_{>t}$ by removing all the edges that take part in a directed cycle of weight 0. Since according to Lemma 3.1, $R_{>t}$ has no negative weight cycles, all the cycles in $\widetilde{R}$ are of strictly positive weight. By the definition of Δ, every simple cycle has a weight of at least Δ. ${R'}_{>t}$ is constructed by taking $\widetilde{R}$ and decreasing all the edge weights by $\epsilon = \frac{\Delta}{|I_{>t}|+1}$. Therefore, the weight of

complete bipartite graph, every maximum weight matching can be extended to produce a matching with the same weight which matches all of the vertices in $I_{>t}$.

$$\forall b \in I_{>t} \quad p(b) \geq 0 \tag{1}$$

$$\forall \langle b_1, b_2 \rangle \in E'_{>t} \quad p(b_1) - p(b_2) < W_{>t}(\langle b_1, b_2 \rangle) \tag{2}$$

$$\forall b \in I_{>t} : v_{>t}(b) > 0 \quad p(b) < v_{>t}(b) \tag{3}$$

Fig. 1. The set of equations that ensures every greedy agent would choose an edge of some maximum weight matching.

Price-Items
Input: A directed graph $G = (I, E, W)$ where all cycles are strictly positive.
Output: a pricing function $p : I \to \Re_{\geq 0}$ such that $p(b') - p(b) \geq -W(\langle b, b' \rangle)$ for every $\langle b, b' \rangle \in E$.

(1) Add a dummy node dum, and draw an edge of weight 0 from dum to every other node.
(2) Compute the shortest path from dum to all nodes of G (there are no negative cycles in G). For every $b \in I$, let $\mathrm{dist}_{dum}(b)$ denote the length of the shortest path from dum to b.
(3) For every item $b \in I$, set $p(b) = -\mathrm{dist}_{dum}(b)$.

Fig. 2. Pricing algorithm.

every simple cycle in $\widetilde{R}$ could have decreased by no more than $|I_{>t}|\,\epsilon < \Delta$, which means that all the cycles in $R'_{>t}$ are of strictly positive weight. ■

[Running example] In Figure 3(b), the thick edges form a directed cycle of weight 0. We remove these edges and subtract ϵ from every remaining edge. We then run Algorithm Price-Items on the obtained graph, which gives the prices presented in red next to each item in Figure 3(b). In this case, the only negative edge (after removing the cycle of length 0) is the edge $\langle d, a \rangle$, whose price is set to $-W'(\langle d, a \rangle) = -(-1-\epsilon) = 1+\epsilon$. Since all the other shortest paths are positive, prices of other items do not change (recall the new price is the negation of the shortest path from the dummy item). When Alice arrives, she picks the unique item in her demand set — item b. Similarly, graphs $G_{>t}, R_{>t}$ of all iterations $t = 0, 1, 2, 3$ are demonstrated in Figure 3(c)-(h).

Consider a directed edge $\langle b_1, b_2 \rangle$ and some cycle it belongs to. The edge $\langle b_1, b_2 \rangle$ came about because we choose a maximal matching where item b_1 was assigned to some buyer a, whereas b_2 was not. If all such cycles have strictly positive total weight, then the edge weights, and the associated prices computed via Price-Items, ensure that agent a prefers b_1 to b_2, effectively removing choices for "wrong" tie breaking. Contrawise, if the edge $\langle b_1, b_2 \rangle$ does belong to some cycle of total weight zero, this implies that the maximum matching is not unique. Ergo, whenever some item along this cycle is first chosen, it is still possible to extend the matching to a maximal weight matching. This is exactly where the dynamic pricing creeps in, subsequent to this symmetry breaking, new prices have to be computed to avoid wrong tie breaking decisions.

We now show that setting prices that satisfy the constraints in Figure 1 ensures that after the arrival of all agents, the social welfare achieved is maximized.

THEOREM 3.1. *A dynamic pricing scheme which calculates prices satisfying the constraints presented in Figure 1 achieves optimal social welfare (a maximum weight matching of G).*

Proof: Recall that M_t is the matching which results from the first $t \in \{0, 1, \dots, |N|\}$ agents taking an item which maximizes their utility and that $G_{>t}$ is the graph of the remaining agents and items after the first t agents arrived and purchased some items. Let $M_{>t}$ be a maximum weight matching of $G_{>t}$, where $M_{>0}$ is a matching that maximizes the social welfare of all the agents, and $M_{>|N|} = \emptyset$. We prove by induction that for every $t \in \{0, 1, \dots, |N|\}$, $\mathsf{SW}(M_t) + \mathsf{SW}(M_{>t}) = OPT$. It follows that the matching $M_{|N|}$ yields optimal social welfare.

For $t = 0$, this claim trivially holds since $\mathsf{SW}(M_{>0}) = OPT$. Assume that for some $t-1$, $\mathsf{SW}(M_{t-1}) + \mathsf{SW}(M_{>t-1}) = OPT$. Let $M_{>t}$ be the maximum weight matching we compute at step 1. of the pricing scheme. When agent t arrives, consider the following cases:

— Agent t does not take any item. From the constraints of type (3), the only case where an agent has no positive utility from any item is if she is matched to an item in M_{t-1} with an edge of weight 0. In this case, $\mathsf{SW}(M_t) = \mathsf{SW}(M_{t-1})$, and by taking $M_{>t}$ to be the same matching as $M_{>t-1}$ without the edge the t-th agent is matched to, $\mathsf{SW}(M_{>t}) = \mathsf{SW}(M_{>t-1})$. We get that $\mathsf{SW}(M_t) + \mathsf{SW}(M_{>t}) = \mathsf{SW}(M_{t-1}) + \mathsf{SW}(M_{>t-1}) = OPT$.

— Agent t takes the item which she is matched to in $M_{>t-1}$. Let v be the value of the t-th agent for the item. Clearly, $\mathsf{SW}(M_t) = \mathsf{SW}(M_{t-1}) + v$. By taking $M_{>t}$ to be the same matching as $M_{>t-1}$ without the edge the t-th agent is matched to, we get $\mathsf{SW}(M_{>t}) = \mathsf{SW}(M_{>t-1}) - v$. We get that $\mathsf{SW}(M_t) + \mathsf{SW}(M_{>t}) = \mathsf{SW}(M_{t-1}) + v + \mathsf{SW}(M_{>t-1}) - v = OPT$.

— Agent $t = a \in N_{>t-1}$ takes an item $b' \in I_{>t-1}$ which is different than $b \in I_{>t-1}$, the item which she is matched to in $M_{>t-1}$. Therefore,

$$v_a(b') - p_{t-1}(b') \geq v_a(b) - p_{t-1}(b). \quad (4)$$

Let $\langle b, b' \rangle \in E_{>t-1}$ be the directed edge from b to b' in $R_{>t-1}$. Its weight $W_{>t-1}(\langle b, b' \rangle) = v_a(b) - v_a(b')$. If $\langle b, b' \rangle$ would have been in $R'_{>t-1}$, then according to constraint (2), we would have had that $p_{t-1}(b) - p_{t-1}(b') < W_{>t-1}(\langle b, b' \rangle) = v_a(b) - v_a(b')$. Rearranging gives us $v_a(b') - p_{t-1}(b') < v_a(b) - p_{t-1}(b)$, which contradicts (4). Therefore, $\langle b, b' \rangle$ was removed from $R'{}_{>t-1}$, which can only happen if the edge is part of a directed cycle of weight 0 in $R_{>t-1}$.

Let $b_1 = b_{\ell+1} = b$, $b_2 = b'$ and let $\langle b_1, b_2 \rangle, \langle b_2, b_3 \rangle, \dots, \langle b_{\ell-1}, b_\ell \rangle, \langle b_\ell, b_{\ell+1} = b_1 \rangle$ be a simple directed cycle of length ℓ and weight 0 in $R_{>t-1}$ in which $\langle b, b' \rangle$ takes part. This cycle corresponds to a cycle of alternating edges in $G_{>t-1}$,

$$(b_1 = b, a_1 = a)\,(a_1, b_2 = b'),(b_2, a_2)\dots(a_{\ell-1}, b_\ell),(b_\ell, a_\ell),(a_\ell, b_{\ell+1} = b_1),$$

where

$$(b_j, a_j) \in M_{>t-1} \text{ and } (a_j, b_{j+1}) \notin M_{>t-1} \quad \text{for every } j \in \{1, \dots, \ell\}.$$

Since the directed cycle is of weight 0, we get that

$$\sum_{j=1}^{\ell} W_{>t}(\langle b_j, b_{j+1} \rangle) = \sum_{j=1}^{\ell} \big(v_{a_j}(b_j) - v_{a_j}(b_{j+1})\big) = 0,$$

which means that the value of the unmatched edges in the directed cycle, $\sum_{j=1}^{\ell} v_{a_j}(b_{j+1})$, is equal to the value of the matched edges, $\sum_{j=1}^{\ell} v_{a_j}(b_j)$.

Let $\widetilde{M}_{>t-1}$ be the matching which is a result of taking $M_{>t-1}$, removing the edges in the set $\{(a_j, b_j)\}_{j \in \{1,\dots,\ell\}}$, and adding the edges of $\{(b_{j+1}, a_j)\}_{j \in \{1,\dots,\ell\}}$; Note that $(a, b') = (a_1, b_2) \in \widetilde{M}_{>t-1}$. Since the edges we added to $\widetilde{M}_{>t-1}$ are of the same value as the edges we removed,

$$\mathsf{SW}(\widetilde{M}_{>t-1}) + \mathsf{SW}(M_{t-1}) = \mathsf{SW}(M_{>t-1}) + \mathsf{SW}(M_{t-1}) = OPT.$$

We define $M_{>t}$ to be a matching comprised of the same edges as $\widetilde{M}_{>t-1}$ except (a, b'). Therefore, $\mathsf{SW}(M_{>t}) = \mathsf{SW}(\widetilde{M}_{>t-1}) - v_a(b')$. Clearly, we have that $\mathsf{SW}(M_t) = \mathsf{SW}(M_{t-1}) + v_a(b')$. We get that $\mathsf{SW}(M_{>t}) + \mathsf{SW}(M_t) = \mathsf{SW}(\widetilde{M}_{>t-1}) - v_a(b') + \mathsf{SW}(M_{t-1}) + v_a(b') = OPT$. This completes the proof of the induction and the theorem.

■

It remains to show that Price-Items satisfies all the constraints in Figure 1. First, we observe that constraints of type (1) are trivially satisfied.

OBSERVATION 3.2. *Price-Items computes prices which satisfy constraints of type* (1).

Proof: This follows since the length of the shortest path from dum to every node is at most the length of the direct edge from dum to this node, *i.e.*, 0. ■

The following property is helpful in proving that constraints of type (2) are satisfied:

LEMMA 3.3. *Let* $G = (I, E, W)$ *be the input graph of Price-Items and let* $p : I \to \Re_{\geq 0}$ *be its output. For every* $\langle b_1, b_2 \rangle \in E$ *we have that* $p(b_2) - p(b_1) \geq -W(\langle b_1, b_2 \rangle)$.

Proof: Since the shortest path from dum to b_2 is no longer than the shortest path from dum to b_1 *plus* the direct edge from b_1 to b_2, we have that

$$\mathrm{dist}_{dum}(b_2) \leq \mathrm{dist}_{dum}(b_1) + W(\langle b_1, b_2 \rangle).$$

Rearranging gives

$$p(b_2) - p(b_1) = -\mathrm{dist}_{dum}(b_2) + \mathrm{dist}_{dum}(b_1) \geq -W(\langle b_1, b_2 \rangle)$$

as desired. ■

We can now establish that constraints of type (2) hold.

LEMMA 3.4. *Price-Items computes prices which satisfy constraints of type* (2).

Proof: By Lemma 3.3, we get that for a given $\langle b_1, b_2 \rangle \in E'_{>t}$,

$$p(b_2) - p(b_1) \geq -W'_{>t}(\langle b_1, b_2 \rangle) = -(W_{>t}(\langle b_1, b_2 \rangle) - \epsilon).$$

Therefore,

$$p(b_1) - p(b_2) \leq W_{>t}(\langle b_1, b_2 \rangle) - \epsilon < W_{>t}(\langle b_1, b_2 \rangle),$$

as desired. ■

For any two items $b, b' \in I$, let $\mathrm{dist}(b, b')$ denote the length of the shortest path from b to b' in $R'_{>t}$. For establishing that constraints of type (3) are met by the prices $p(b)$'s computed by Price-Items, we need the following lemma.

LEMMA 3.5. *Let* b_ℓ *be some vertex with* $p(b_\ell) > 0$*, and let* $dum, b_0, b_1, \ldots, b_\ell$ *be a shortest path from the dummy node* dum *to* b_ℓ*. For every* $i \in \{0, 1, \ldots, \ell\}$*,* $p(b_i) = -\mathrm{dist}(b_0, b_i)$.

Proof: Let b_i a vertex on the shortest path from b_0 to b_ℓ. Since every sub-path of a shortest path is also a shortest path, it must be that $dum, b_0, \ldots, b_i$ is a shortest path from dum to b_i, and that $\mathrm{dist}_{dum}(b_i) = W(\langle dum, b_0 \rangle) + \mathrm{dist}(b_0, b_i) = \mathrm{dist}(b_0, b_i)$. Therefore, $p(b_i) = -\mathrm{dist}_{dum}(b_i) = -\mathrm{dist}(b_0, b_i)$ as desired. ■

We get the the following two corollaries.

COROLLARY 3.3. $p(b_0) = 0$.

COROLLARY 3.4. *For every* $i \in \{0, 1, \ldots, \ell - 1\}$, $p(b_i) - p(b_{i+1}) = W_{>t}(\langle b_i, b_{i+1}\rangle) - \epsilon$.

Proof: Since every sub-path of a shortest path is also a shortest path, we get that $\mathrm{dist}(b_0, b_{i+1}) = \mathrm{dist}(b_0, b_i) + W'_{>t}(\langle b_i, b_{i+1}\rangle)$. From Lemma 3.5, we get that $p(b_i) = -\mathrm{dist}(b_0, b_i)$ and

$$\begin{aligned} p(b_{i+1}) &= -\mathrm{dist}(b_0, b_{i+1}) \\ &= -\mathrm{dist}(b_0, b_i) - W'_{>t}(\langle b_i, b_{i+1}\rangle) \\ &= p(b_i) - (W_{>t}(\langle b_i, b_{i+1}\rangle) - \epsilon), \end{aligned}$$

where the last equality follows by the definition of $W'_{>t}$. ■

We now prove that all the constraints of type (3) are met.

LEMMA 3.6. *For every* $b \in I_{>t}$ *which is matched in* $M_{>t}$ *by a non-zero weight edge,* $p(b) < v_{>t}(b)$.

Proof: Assume for the purpose of reaching a contradiction that there exists some $b = b_\ell$ which is matched in M_t via an edge of strictly positive weight for which $p(b) \geq v_{>t}(b)$. Let $dum, b_0, b_1, \ldots, b_\ell$ be a shortest path from dum to b_ℓ in the graph processed in Price-Items. According to Corollary 3.4, for every $i \in \{0, 1, \ldots, \ell - 1\}$, $p(b_i) - p(b_{i+1}) = W_{>t}(\langle b_i, b_{i+1}\rangle) - \epsilon$. Summing over all i's gives us

$$\sum_{i=0}^{\ell-1} W_{>t}(\langle b_i, b_{i+1}\rangle) = p(b_0) - p(b_\ell) + \ell\epsilon < -p(b) + \Delta, \tag{5}$$

where the inequality stems from the fact that $p(b_0) = 0$ (Corollary 3.3), $b_\ell = b$, $\ell < |I_{>t}|$ and $\epsilon = \frac{\Delta}{|I_{>t}|+1}$. Let a be the vertex that b is matched to in $M_{>t}$. According to the definitions of the weights of edges in $R_{>t}$, we get that the weight of the edge $\langle b, b_0\rangle \in E_t$ in $R_{>t}$ is

$$W_{>t}(\langle b_\ell, b_0\rangle) = v_a(b) - v_a(b_0) \leq v_{>t}(b) \leq p(b), \tag{6}$$

where the first inequality is due to the definition of $v_{>t}(b)$, and the second inequality is due to our initial assumption. Combining (5) with (6) yields that the weight of the cycle $\langle b_0, b_1\rangle, \langle b_1, b_2\rangle, \ldots, \langle b_{\ell-1}, b_\ell\rangle, \langle b_\ell, b_0\rangle$ in $R_{>t}$ is $\sum_{i=0}^{\ell} W_{>t}(\langle b_i, b_{i+1 \mod \ell+1}\rangle) < \Delta$. Since Δ is the minimal weight of a positive cycle in $R_{>t}$, we get that either the weight of the cycle is negative, which contradicts Lemma 3.1, or the cycle is of weight 0, contradicting the fact the we delete every edge that takes part in some cycle of weight 0 in $R_{>t}$ from ${R'}_{>t}$. ■

4. NO OPTIMAL DYNAMIC PRICING SCHEME FOR COVERAGE VALUATIONS

We show an instance with agents with coverage valuations[3] for which no dynamic pricing scheme guarantees an optimal allocation. Interestingly, this instance admits Walrasian prices and has a unique optimal allocation, so no combination of these conditions is sufficient to imply optimal dynamic pricing schemes.

THEOREM 4.1. *There exists an instance with agents with coverage valuations such that no dynamic pricing scheme guarantees more than a fraction* $\frac{7.5}{8}$ *of the optimal social welfare. This instance admits Walrasian prices.*

[3]The class of coverage valuations is a strict subclass of submodular valuations.

Proof: Let $I = \{a, b, c, d\}$ be a set of items and $N = \{1, 2, 3, 4\}$ be a set of agents. Agents $2, 3$ and 4 are unit demand with the following valuation functions:

$$v_2(S) = \begin{cases} 2 & S \cap \{a, b\} \neq \emptyset \\ 0 & otherwise \end{cases}, v_3(S) = \begin{cases} 2 & S \cap \{a, c\} \neq \emptyset \\ 0 & otherwise \end{cases}, v_4(S) = \begin{cases} 1 & S \cap \{d\} \neq \emptyset \\ 0 & otherwise \end{cases}.$$

In addition, agent 1 has the following coverage valuation:

$$v_1(S) = \begin{cases} 2 & S = \{b\}, S = \{c\} \\ 3 & S = \{a\}, S = \{d\} \\ 3.5 & S = \{a, b\}, S = \{a, c\}, S = \{d, b\}, S = \{d, c\}, S = \{a, d\} \\ 3.75 & S = \{a, b, d\}, S = \{a, c, d\} \\ 4 & \{b, c\} \subseteq S \end{cases}.$$

Coverage valuation: To see that this is a coverage valuation, consider the following explicit representation. Let $\{e_1, e_2, e_3, e_4, e_5, e_6, e_7, e_8\}$ be the set of elements, with weights $w(e_1) = w(e_5) = 5/4$ and $w(e_i) = 1/4$ for $i \neq 1, 5$. Item a covers the set $\{e_1, e_2, e_5, e_6\}$, item b covers the set $\{e_1, e_2, e_3, e_4\}$, item c covers the set $\{e_5, e_6, e_7, e_8\}$, and item d covers the set $\{e_1, e_4, e_5, e_8\}$.

Unique optimal allocation: The unique optimal allocation is to allocate item a to agent 1, item b to agent 2, item c to agent 3 and item d to agent 4. This allocation obtains social welfare of 8.

Walrasian prices: One can easily verify that the unique optimal allocation along with pricing each item at 1 is a Walrasian equilibrium.

We now show that no dynamic pricing scheme guarantees more than a fraction $\frac{7.5}{8}$ of the optimal allocation. In order to guarantee an optimal allocation, the following conditions must be satisfied:

— Agent 4's utility from item d should be strictly positive; i.e.,

$$p(d) < v_4(d) = 1. \quad (7)$$

— Agent 1 should strictly prefer item a over item d, i.e.,

$$v_1(a) - p(a) > v_1(d) - p(d) \Rightarrow p(a) < p(d). \quad (8)$$

— Agent 2 should strictly prefer item b over item a, i.e.,

$$v_2(b) - p(b) > v_2(a) - p(a) \Rightarrow p(b) < p(a). \quad (9)$$

— Agent 3 should strictly prefer item c over item a, i.e.,:

$$p(c) < p(a). \quad (10)$$

— Agent 1 should strictly prefer item a over the bundle $\{b, c\}$, i.e.,

$$v_1(a) - p(a) > v_1(\{b, c\}) - p(b) - p(c) \Rightarrow p(b) + p(c) - p(a) > 1. \quad (11)$$

Combining Equations (7) and (8) implies that $p(a) < 1$, while combining Equations (9), (10) and (11) yields $p(a) > 1$. Therefore, for every prices one might set, the adversary can set an order for which the first agent picks a different item than the one allocated to her in the optimal allocation.

Remark: note that the valuation function of agent 1 is not gross substitutes. In particular, her demand under prices $p(a) = p(c) = p(d) = 0$ and $p(b) = \epsilon$ is $\{b, c\}$, but if the price of item c increases to 2, then the unique bundle in the demand of agent 1 is $\{a, d\}$. ■

5. A 1/2-APPROXIMATE STATIC PRICING SCHEME FOR ANY CLASS OF VALUATIONS

In this section we show that, given a partition of the items into bundles, pricing each bundle half of its value to the buyer guarantees half of the social welfare of the partition. Let $\mathcal{B} = \{B_1, B_2, \ldots, B_n\} \in \left(2^I\right)^n$ be a partition of the items such that $\bigcup_i B_i = I$ and for every $i \neq j$ $B_i \cap B_j = \emptyset$. Let $W = \sum_i v_i(B_i)$. We have the following:

THEOREM 5.1. *Let $p : \mathcal{B} \to \Re_{\geq 0}$ be static bundle prices such that for every i, $p(B_i) = v_i(B_i)/2$. This pricing scheme achieves a welfare of at least $W/2$.*

Proof: Let x be an allocation which is a result of agents arriving at an arbitrary order, each taking their favorite bundles. Notice that the utility of an agent for acquiring the bundles in x_i is $u_i(x_i, P) = v_i(\bigcup_{B \in x_i} B) - \sum_{B \in x_i} p(B)$. Let $\mathbb{I}_i$ be an indicator which gets 1 if bundle B_i was acquired by some agent and 0 otherwise. Rearranging and summing over all the agents gives us:

$$\sum_i v_i\left(\bigcup_{B \in x_i} B\right) = \sum_i \left(u_i(x_i, P) + \sum_{B \in x_i} p(B)\right)$$
$$= \sum_i u_i(x_i, P) + \mathbb{I}_i p(B_i). \tag{12}$$

We show that for every i, $u_i(x_i, P) + \mathbb{I}_i p(B_i) \geq v_i(B_i)/2$. Using (12) this is enough to prove the claim. For some i, either bundle B_i is purchased by some agent, in which case $\mathbb{I}_i p(B_i) = v_i(B_i)/2$. Otherwise, when agent i arrived, she could have purchased bundle B_i, for which she would have gotten a utility of $v_i(B_i) - p(B_i) = v_i(B_i)/2$. Since she bought the bundles which maximized her utility, her utility can only be greater than that, meaning $u_i(x_i, p) \geq v_i(B_i)/2$. ■

6. OPTIMAL PRICING SCHEMES FOR GROSS SUBSTITUTES VALUATIONS WITH UNIQUE OPTIMAL ALLOCATIONS

In Section 4 we give an example where there is a unique optimum, there exist Walrasian prices over the items, and no dynamic bundle pricing scheme can guarantee an optimal outcome.

Here, we show that in the case of Gross Substitute valuations, a unique optimum implies the existence of *static prices* that guarantee an optimal allocation (for any order of arrival).

We first show that in case of a unique optimum, an optimal dynamic bundle-pricing scheme implies an optimal static bundle-pricing scheme:

OBSERVATION 6.1. *Let $\mathbf{v} = (v_1, \ldots, v_n)$, where $v_i : 2^I \mapsto \Re_{\geq 0}$, and let $\langle \mathbf{v}, I \rangle$ be an instance where $\mathcal{B} = \{B_1, \ldots, B_n\}$ is the unique partition of items that maximizes social welfare. If there exists an optimal dynamic bundle-pricing scheme, then there must exist an optimal static bundle-pricing scheme.*

Proof: Let $p_1 : \mathcal{B} \to \Re_{\geq 0}$ be the initial prices the optimal dynamic pricing scheme gives to the bundles. We claim that sticking to these prices throughout the process guarantees an optimal allocation as well. Without loss of generality, assume that agents with lower index arrive earlier and that the i-th agent to arrive is the first agent whose choice $X \neq \{B_i\}$ (could be that $X = \{B_j\}$, $j \neq i$, could be that $x = \{B_i, B_j, \ldots\}$, $j \neq i$, and could be that $X = \emptyset$).

It must be the case that $u_i(p_1, X) \geq u_i(p_1, B_i)$. Therefore, if this agent arrives first, she is not guaranteed to take $\{B_i\}$ since this not the unique bundle that maximizes her utility. This contradicts the optimality of the dynamic pricing scheme. ■

We next show how to compute a static pricing scheme for the case of gross substitutes with unique optimum. We do so via a combinatorial algorithm inspired by Murota [Murota 1996a,b].[4]

Given some set of items $A \subseteq I$, we define the sets of items local to A as following $\mathbf{Local}(A) = \{B \neq A \subseteq I : |B \setminus A| \leq 1 \text{ and } |A \setminus B| \leq 1\}$ We present the following alternative definition of gross substitute valuations [Gul and Stacchetti 1999]:

DEFINITION 6.1. *A valuation $v : 2^I \to \Re_{\geq 0}$ satisfies the gross substitute condition whenever the following holds: for every item prices $p : I \to \Re_{\geq 0}$, for every $A \subseteq I$ such that $A \notin D(I, p)$ there exists $B \in \textbf{\textit{Local}}(A)$ such that $u(B, p) > u(A, p)$.*

We refer to this characterization as the *local improvement* property (LI).

Given a set of gross-substitute valuations and items $\langle \mathbf{v}, I \rangle$, let $\mathcal{B} = \{B_1, \ldots, B_n\}$ be the unique optimal allocation. We compute the prices $p : I \to \Re_{\geq 0}$ as follows:

(1) Let $D = \{d_1, \ldots, d_n\}$ be a set of dummy items (one for each agent), $I' = I \cup D$ be the set of items after we added the dummy items. We extend every valuation function v_i to the domain $2^{I'}$, where $v_i(X) = v_i(X \cap I)$ (i.e., the dummy items have no effect on the value of the bundle). Define $\mathcal{B}' = \{B'_1, \ldots, B'_n\}$ where every bundle $B'_i = \{B_i \cup \{d_i\}\}$ receives an additional dummy item.
(2) Let $R = \langle V = I', E \subset V \times V, W : E \to \Re \rangle$ (the exchange graph) be a weighted directed graph where:
   — $E = \{\langle a, b \rangle \in I'^2 : a \in B'_i, b \in I' \setminus B'_i \text{ for every } i\} \setminus D^2$: I.e., there is an edge from every item in some bundle B'_i to every item not in B'_i, *unless* the two items are dummy items.
   — Let $e = \langle a, b \rangle$ where $a \in B'_i$ of some agent i be an edge in the graph. $W(e) = v_i(B'_i) - v_i(B'_i - a + b)$, i.e., the value of the agent from bundle B'_i *minus* the value she gets if she exchanges item a for item b.
(3) Let $\delta > 0$ be the weight of a minimum weight cycle in R (in the full version we show that all the cycles in R are of strictly positive weight). Let $\gamma > 0$ be the weight of the minimum weight path out of all the paths from any vertex to any dummy vertex (in the full version we show that all such paths are of strictly positive weight). Let $\epsilon = \frac{\min\{\delta, \gamma\}}{n+1}$.
(4) Update the weights by setting $W(e) \leftarrow W(e) - \epsilon$ for every edge e in the graph.
(5) Price the items using algorithm Price-Items (Figure 2) with graph R as input.

In the full version, we prove the following theorem[5]:

THEOREM 6.2. *Item prices p computed above achieve optimal welfare irrespective of the order of arrival of the buyers.*

7. MODELS OF UNCERTAINTY

In this section we generalize our results beyond the full information setting. We consider two models of uncertainty and in both cases we give dynamic pricing that achieves social welfare close to optimum, up to the inherent uncertainty in the models. Notably, all of these elements of uncertainty can be combined, and our techniques will give social welfare arbitrarily close to optimal.

[4]See [Paes Leme 2014] for a concise description on how Murota's work relates to the computation of Walrasian prices for GS valuations.

[5]Independently, Paes Leme and Wong [2015] defined robust Walrasian pricing where there is no overlap between the demand sets of different buyers, and showed that for Gross substitute valuations with unique optima, such prices exist. Viewed from our perspective, this gives static prices that achieve optimal welfare for any order of arrival.

7.1. Mixed Bayesian/Worst case setting

We now consider a model where the number of buyers per day is known, the different types of buyers are known, and the fraction of buyers from each type is approximately known. What is not known is when the different buyers will arrive and one allows an adversary to determine this order. As in the full information setting, making such a worst case assumption about the order of arrival can be efficiently dealt with by setting dynamic prices.

The mixed model is as follows:

— There is a set of k user types, $t_1, \ldots, t_k$, each of which is a unit demand valuation.
— We assume that the entire population will eventually arrive, and that the population size is large. For simplicity we also assume that the population size is known to seller, although this could be extened to be known to within some small error.
— The population of buyers is generated by repeatedly and independently sampling from the type distribution $\mathcal{D}$. The probability of any type being chosen is bounded from bellow by $\omega(1/n)$.
— The order of arrival of individuals from the population is determined by an adversary, this can be a function of the realized valuations of the entire population.

For this setting we prove the following theorem. The proof is deferred to the full version.

THEOREM 7.1. *One can compute dynamic prices for the mixed model that achieves $1-\epsilon$ approximation of the optimal social welfare with probability of at least $1-\epsilon$.*

7.2. Fuzzy valuations

Given a valuation v, we define the set of ϵ-approximate valuation functions for v

$$\widetilde{V}(\epsilon, v) = \{v' : \text{ for all items } b \ |v'(b) - v(b)| < \epsilon\}.$$

We model bounded rationality agents, and sellers with fuzzy knowledge of the valuations. For every i, let $\widetilde{V}_i = \widetilde{V}(\epsilon, v_i)$. Our model is as follows:

— For every agent i there is some central valuation function v_i.
— Whenever computing the demand set for buyer i, we assume that her valuation is some arbitrary valuation function $\widetilde{v}_i \in \widetilde{V}_i$.
— The seller does not know v_i, but knows ϵ and some arbitrary $\widetilde{w}_i \in \widetilde{V}_i$.

For this setting we prove the following theorem. The proof is deferred to the full version.

THEOREM 7.2. *For any $\epsilon > 0$, one can compute dynamic prices for the fuzzy valuations model that achieve a social welfare of at least $OPT - \epsilon$.*

ACKNOWLEDGMENTS

We are grateful to Claire Mathieu and Orit Raz for helpful discussions.

Bibliography

Liad Blumrosen and Thomas Holenstein. 2008. Posted prices vs. negotiations: an asymptotic analysis. In *EC*. 49.

Shuchi Chawla, Jason D. Hartline, and Robert D. Kleinberg. 2007. Algorithmic pricing via virtual valuations. In *ACM Conference on Electronic Commerce*. 243–251.

Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. 2010a. Multi-parameter mechanism design and sequential posted pricing. In *STOC*. 311–320.

Shuchi Chawla, David L. Malec, and Balasubramanian Sivan. 2010b. The Power of Randomness in Bayesian Optimal Mechanism Design. In *the 11th ACM Conference on Electronic Commerce (EC)*.

Ilan Reuven Cohen, Alon Eden, Amos Fiat, and Lukasz Jeż. 2015. Pricing online decisions: Beyond auctions. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on discrete algorithms*. SIAM, 73–91.

Vincent P. Crawford and Elsie Marie Knoer. 1981. Job matching with heterogenous firms and workers. *Econometrica: Journal of the Econometric Society* 49 (1981). Issue 2.

Michal Feldman, Nick Gravin, and Brendan Lucier. 2013. Combinatorial walrasian equilibrium. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*. ACM, 61–70.

Michal Feldman, Nick Gravin, and Brendan Lucier. 2015. Combinatorial auctions via posted prices. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 123–135.

Faruk Gul and Ennio Stacchetti. 1999. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory* 87, 1 (1999), 95–124.

Justin Hsu, Jamie Morgenstern, Ryan M. Rogers, Aaron Roth, and Rakesh Vohra. 2016. Do Prices Coordinate Markets?. In *Proceedings of the forty-eight annual ACM symposium on Theory of computing*.

Richard M Karp, Umesh V Vazirani, and Vijay V Vazirani. 1990. An optimal algorithm for on-line bipartite matching. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing*. ACM, 352–358.

Alexander S Kelso Jr and Vincent P Crawford. 1982. Job matching, coalition formation, and gross substitutes. *Econometrica: Journal of the Econometric Society* (1982), 1483–1504.

Herman B Leonard. 1983. Elicitation of Honest Preferences for the Assignment of Individuals to Positions. *Journal of Political Economy* 91, 3 (1983), 461–79. http://EconPapers.repec.org/RePEc:ucp:jpolec:v:91:y:1983:i:3:p:461-79

Kazuo Murota. 1996a. Valuated matroid intersection I: Optimality criteria. *SIAM Journal on Discrete Mathematics* 9, 4 (1996), 545–561.

Kazuo Murota. 1996b. Valuated matroid intersection II: Algorithms. *SIAM Journal on Discrete Mathematics* 9, 4 (1996), 562–576.

Noam Nisan and Ilya Segal. 2006. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory* 129 (2006), 192–224.

Renato Paes Leme. 2014. Gross substitutability: An algorithmic survey. *preprint* (2014).

Renato Paes Leme and Sam Chiu-wai Wong. 2015. Computing Walrasian Equilibria: Fast Algorithms and Economic Insights. *CoRR* abs/1511.04032 (2015). http://arxiv.org/abs/1511.04032

L. Walras. 1874. *Éléments d'économie politique pure; ou, Théorie de la richesse sociale*. Corbaz. https://books.google.com/books?id=crUEAAAAMAAJ

A. AN ILLUSTRATION OF OUR RUNNING EXAMPLE

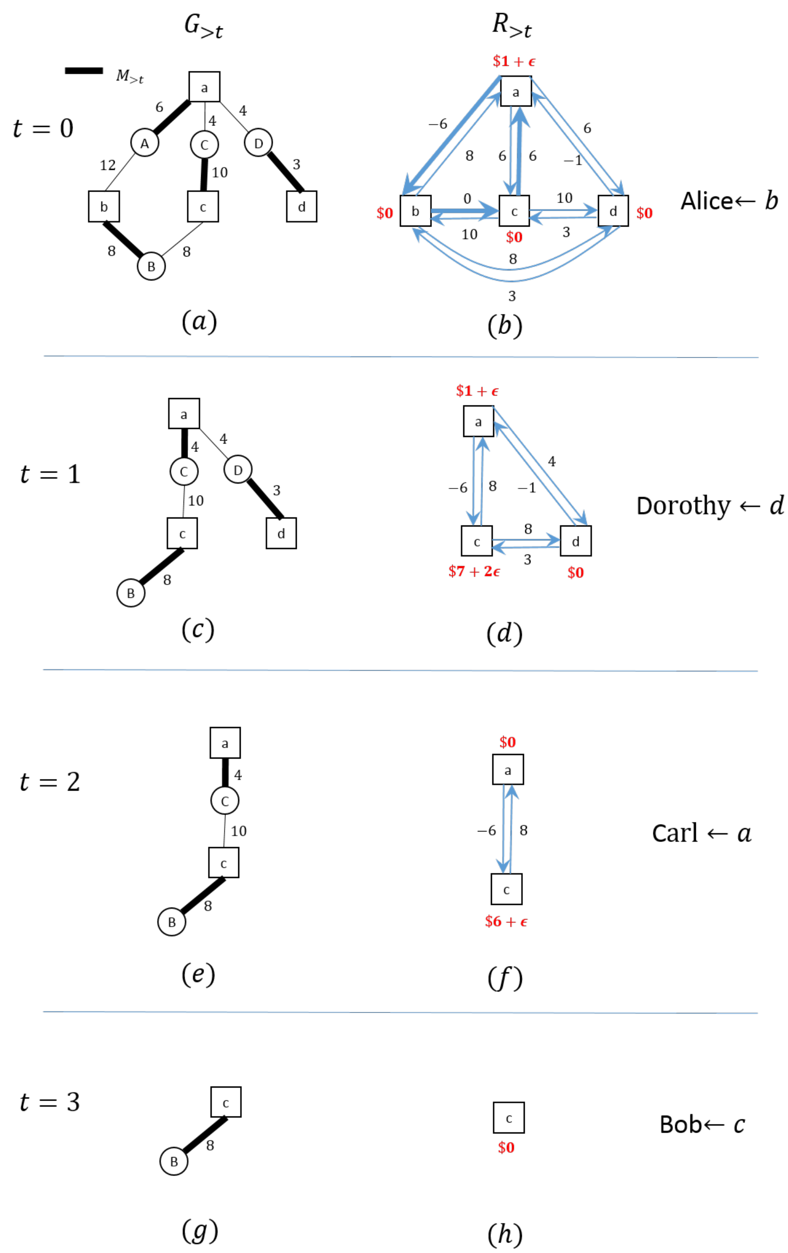

Fig. 3. Phases $t = 0, 1, 2, 3$ of our running example. Squares represent items and circles represent buyers. Every row represents a phase in the process, where a single buyer arrives. On the left one sees the graph representing the valuations of the remaining buyers and items (graphs labeled (a), (b), (c) and (d), where thick edges represent a maximum weight matching in the graph. Graphs labeled (b), (d), (f) and (h) give the graphs $R_{>t}$ from which the dynamic are computed. Directed cycles of length 0 (if any) are represented by thick edges, after they are discarded, prices are computed via Algorithm Price-Items. On the very right one sees the next buyer to arrive as well as the item she chooses (based upon the pricing, and breaking ties for the sake of this example.